REMARKS

Reconsideration and allowance of this application are respectfully requested.

By this Amendment, claim 64 has been cancelled without prejudice or disclaimer of its subject matter, claims 41, 45, 48, 49, 52, 55, 56, and 59-63 have been amended and claims 70-86 have been added. The title of the invention has been amended and a new *Abstract Of The Disclosure* has been provided. No new matter has been added.

Applicant thanks Primary Examiner Geckil for the courtesies extended its representative, Dr. Brian Siritzky, during various telephone interviews and during the personal interview of September 20, 2001. During that interview the Examiner and Dr. Siritzky discussed all of the claims and the status of applicant's interference request.

REJECTIONS UNDER 35 U.S.C. § 112, SECOND PARAGRAPH

The Examiner rejected claims 49, 50, 52, 57-59, 62-63, 67, and 68 under 35 U.S.C. § 112, second paragraph, as being indefinite. The claims have been amended to deal with the issues raised by the Examiner.

Specifically, claim 52 has been amended to depend from claim 51 and not claim 49. Applicant thanks the Examiner for pointing out this typographical error in the claim. In addition, the Examiner stated that the phrase "the server" in claim 52 lacks antecedent basis. Claim 52, as amended, depends from claim 51 which recites, *inter alia*: "resolving the ARL to identify a server in the domain." (Emphasis added). Claim 52, as amended, recites "The method as described in claim 51 wherein the identified server is selected from a set of repeater servers based data identifying a requesting user's location and data identifying current costs between a group containing the requesting user and servers in the set of repeater servers."

Claim 59 has been amended as suggested by the Examiner to correct the noted typographical error.

Claims 62 and 67 have been amended to clarify the antecedent basis of the terms noted by the Examiner.

In view of the amendments to the claims, withdrawal of this rejection is respectfully requested.

OBJECTIONS TO THE SPECIFICATION

The Examiner objected to the Specification under 37 CFR 1.71 as failing to provide an enabling disclosure. In particular, the Examiner asserted that the specification did not support the first level name server of claim 45 or the details of the first and second level name servers of claim 41. Claim 41 has been amended to remove the reference to the first and second level name servers, and claim 45 has been amended to remove reference to any name server. Claim 41, as amended, recites "a repeater selector mechanism constructed and adapted to identify, for a particular client machine, an appropriate repeater server from the set of repeater servers" and further, that an "embedded object identified by the modified embedded object URL is served from a given one of the repeater servers as identified by the repeater selector mechanism." Claim 45 has been amended to refer to the repeater selector mechanism.

The Examiner objected to the specification for not supporting claim 48. Claim 48 has been amended in accordance with the Examiner's suggestions.

The Examiner asserted that there is no support in claims 52, 55, 56, 60, 61 and 63 for the term "current Internet traffic conditions." The claims have been amended (or replaced) to clarify that the identified server is selected using data identifying current costs between a group containing the requesting client and a set of repeater servers. Support for these claims is found, e.g., at pages 19-24. Specifically, describing the Best Repeater Selector (BRS), the application states (at page 23, lines 6-21):

> Given a client network address and the three tables described *above [Group Reduction Table, Link Cost Table* and *Load Table*, at page 20]:
> E1. Determine which group the client is in using the Group Reduction Table.
> E2. For each repeater in the Link Cost Table and Load Table, determine that repeater's combined cost as follows:
> E2a. Determine the maximum and current load on the repeater (using the Load Table).
> E2b. Determine the link cost between the repeater and the client's group (using the Link Cost Table).
> E2c. Determine the combined cost * * *
> E3. Select a small set of repeaters with the lowest cost.
> E4. Select a random member from the set.

At step E1, using the requesting client's network address, i.e., a location of the requesting client in the network, the system determines which group the client is in. "**The Group Reduction Table** places every network address into a group. * * * **The Link Cost Table** * * *

specifies the current cost between each repeater and each group. * * * Over time, the table will * * * accurately reflect the relative cost of transmitting a file between the repeater and a member of the group." *Pg. 20*, lines 8-16, bold in original. "The term *Group* . . . refers to an IP 'address group'" *Pg. 24, line 8.* I.e., the system selects a repeater, *inter alia,* based on "the cost between each repeater and each group."

The Examiner rejected claims 41, 42, 45, and 48-68 under 35 U.S.C. § 112, first paragraph, for the reasons set forth in his objections to the specification. In view of the above, withdrawal of the objections to and rejections of the specification are respectfully requested.

PRIOR ART REJECTIONS

The Examiner rejected claims 41, 42, 48-51, 53, 54, 57-59, 62 and 64-69 under 35 U.S.C. § 103 as being unpatentable over Graber, U.S. Patent No. 5,712,979, ("*Graber '979*"). The Examiner also applied Graber, U.S. Patent No. 5,812,769 ("*Graber '769*") to reject claims 51, 57-59, 62, 64, 67 and 68 as being clearly anticipated under 35 U.S.C. § 102(e). While apparently rejecting the claims over two references, Applicant notes that the disclosures of Graber '979 and Graber '769 are identical. Graber '769 is a continuation application of Graber '979, with identical descriptions and drawings, therefore Applicant understands these to be the same rejection.

The Examiner also rejected claims 45, 52, 55, 56, 60, 61 and 63 as unpatentable under 35 U.S.C. § 103 over Graber in view of Bonnaure.

As to claims 41, 42, 48-51, 53, 54, 57-59, 62 and 64-69, the Examiner's position is that Graber '979 taught the invention substantially as claimed. The grounds for this rejection are respectfully traversed.

Graber relates to a method and apparatus for attaching navigational history information to universal resource locators (URLs) on a World Wide Web (WWW) page. *Graber '979,* Abstract. More specifically, Graber '979 relates to "[a] method and apparatus for tracking the navigation path of a user that has been directed to a second site on the WWW from a first site on the WWW." *Graber '979,* Abstract.

In many cases, a WWW user may be able to access a particular web page from any of a number of sources. This is the case, e.g., when the user may subscribe to a service from a number of different locations on the WWW. Graber is trying to solve two related problems

caused by the fact that a particular location on the WWW may have been reached from a number of different paths. First, when a user reaches a particular location on the WWW, it is often useful to know where a user came from, e.g., so that the appropriate party can be credited with sending the user to that location. This is the example used by Graber, where it describes a number of different co-marketers of an on-line service. Second, when a user reaches a particular location on the WWW, it is often useful to know the path the user took to get to that location so that if the user selects to navigate backwards, the correct reverse path is maintained.

> A URL is received at the second WWW site when the user is directed from the first site to the second site. At the second WWW site, information representative of an identity of the first WWW site is captured . . .. A destination web page is determined for the user, and a revised destination web page is formed by attaching a second code representative of the identity of the first WWW site into at least one selected web page link associated with the destination web page. The revised destination web page is then transmitted to the user.
> *Graber '979*, Abstract

Importantly, Grabber teaches modifying a URL so that when a user moves between pages on the same web site, UNIX symbolic link information is retained. As noted above, Grabber is trying to solve the problem of determining where a user came from so that when the user selects a "BACK" option on his browser, the user will be sent back to the appropriate place.

However, there is nothing in Graber to teach or in any way suggest modifying URLs to move content between different servers. For example, Graber teaches that

> . . . , if the user clicks on the 'Enroll on OLS box" on page 514a, special redirecting program (redirect.cgi) is triggered on web server 142 for redirecting the user from the page represented by URL 514 to the OLS enrollment page represented by URL 518.
> *'979*, col. 10, lines 57-61, emphasis added.

That is, using the example URLs provided by Graber, the URLs are modified from:

"WWW.OLS.COM/ . . ./INFO/INFO.P1"

to

"WWW.OLS.COM/CM2/. . ./ENROLL/ENROLL.P1"

So, when the user clicks on the "Enroll on OLS box" in order to enroll in the online service, the redirect.cgi program updates the path information in the URL. Notably, both URLs refer to "WWW.OLS.COM" -- the same server.

Grabber provides a pseudocode version of the *redirect.cgi* program in Table I (col. 11 of Graber '979) which shows that a URL is modified to create a new URL (new_absolute_url = concatenate(new_directory, relative_url)) to which the browser is redirected (redirect_browser(new_absolute_url)). The "*new_directory*" field is determined based on removing levels from the last_url input parameter, therefore the new URL will always go to the same computer as specified by the last_url. In other words, the new URL does not redirect a user to another server, just to another location in the server on which it is run.

> Thus, the redirect.cgi program accepts as arguments the current URL of the user (e.g., URL 514) and a destination URL representing the location to which the user desires to move (e.g., URL 518). The program then strips the ". . ./Info/Info.P1" portion off of the current URL 514, and replaces the striped portion with the ". . . /Enroll/Enroll.Pl" portion of destination URL 518 to form a new URL which is then used for redirecting the user to the page represented by URL 518. The redirect.cgi program is significant to the operation of the present invention because, among other things, this program allows the UNIX symbolic link information that was originally passed when the user arrived at the home page of OLS site 128 to be retained as the user moves between pages at OLS site 128. Thus, the redirecting.cgi program insures that the UNIX symbolic link information provided by a co-marketer will be present when the enrollment means 145 attempts to enroll the user on OLS 140.
>
> The redirect.cgi . . . represents a first preferred system for retaining at site 128 the UNIX symbolic link information (that was originally passed when the user originally arrived at OLS site 128 from a previous site) as the user moves between web pages at OLS site 128.
>
> *Graber '979*, col. 11, lines 31-54, emphasis added.

The Examiner asserts that Graber '979 teaches "a routine for modifying at least one embedded object URL or a link of a web page . . . (col. 10, lines 57-68 and col. 11 line 30 *et seq*)" *Paper No. 23, pg. 10, item 7.* However, to the extent that Graber teaches any form of URL modification, all that Graber modifies is links, not embedded object URLs. The Examiner acknowledges this in his Action (Paper No. 23):

> * * * (see for example column 13 . . . for external URL links being appended and col 14 . . . for the destination page which includes the URLs having the appended codes being passed to the

> user and the user executing or fetching these links." *Paper No. 23*, pg. 11, item 8, emphasis added.

The Examiner points to no teaching or suggestion of modifying embedded object URLs.

In addition, in the rejection of claims 51, 57-59, 62, 64, 67 and 68 as being clearly anticipated by Graber '769 under 35 U.S.C. § 102(e), the Examiner stated that Graber's system

> . . . modified and appended the URL links into the requested page and sent the page with the appended links so that the user will select one of the embedded or appended object URL to fetch the desired object from the destination server identified by the link URL. . . . Claims do not recited automatically fetching of the embedded objects pointed by the appended URL. Thus, in the Grabber [sic] et al system user selects the modified appended URL link by clicking on the link and the embedded object is resolved and received from the identified destination server. This operation of the Grabber [sic] et al system reads on these claims because of the broad recitation of the claim language."
> *Paper No. 23,* pgs. 13-14, item 12.

The Examiner also acknowledged this point in the personal interview of September 20, 2001. There the Examiner agreed that links had to be explicitly selected (clicked on) by a user, whereas other objects might be automatically loaded by the browser. The Examiner's position was that the automatic loading was not recited in the claims. Applicant respectfully disagrees. Claim 41 recites "a routine for modifying at least one embedded object URL of a web page to include a hostname prepended to a domain name and path". Embedded objects are, by definition, automatically downloaded on a user's computer when encountered in an HTML page.

HTML offers a mechanism for embedding generic media objects and applications in HTML documents. E.g., the HTML OBJECT element (together with its more specific ancestor elements IMG and APPLET) provides a mechanism for including images, video, sound, mathematics, specialized applications, and other objects in a document. HTML 4 Specification §2.3.4. (See, e.g., "*http://www.w3.org/TR/html4/intro/intro.html#h-2.2*").

HTML links are not the same as HTML objects. By definition, links require explicit action on the part of users (e.g., clicking on the link) in order to visit a linked to object. Objects, on the other hand, are automatically downloaded on a user's computer when encountered in an HTML page. According the HTML 4 Specification, "a link is a connection from one Web resource to another." Chap. 12 (see, e.g., "*http://www.w3.org/TR/html4/struct/links.html*").

Linked URLs are not automatically loaded when a web page is loaded. See, e.g., the HTML 4.01 Specification § 12.1.1 "Visiting a linked resource" "The default behavior associated with a link is the retrieval of another Web resource. This behavior is commonly and implicitly obtained by selecting the link (e.g., by clicking, through keyboard input, etc.)." On the other hand, as to objects, "HTML's multimedia features allow authors to include images, applets (programs that are automatically downloaded and run on the user's machine), video clips, and other HTML documents in their pages" HTML Specification, § 13.1, emphasis added (see, e.g., "*http://www.w3.org/TR/html4/struct/objects.html*"). Thus, when a browser loads an HTML web page, it automatically loads all the embedded objects in that page. More specifically, when a browser loads an HTML page, it loads the embedded objects associated with each embedded object URL in the page.

The following HTML excerpt contains two links, one whose destination anchor is an HTML document named "chapter2.html" and the other whose destination anchor is a GIF image in the file "forest.gif":

```
<BODY>
...some text...
<P>You'll find a lot more in  <A href="chapter2.html">chapter two</A>.
See also this <A href="../images/forest.gif">map of the enchanted forest.</A>
</BODY>
```

Links are designated by "<A href = "link">. By activating these links (by clicking with the mouse, through keyboard input, voice commands, etc.), users may visit these resources. Note that the "href" attribute in each source anchor specifies the address of the destination anchor with a URI.

The following example shows how to include an object (i.e., something that is automatically loaded) (not a link) into a document (an image embedded object is designated by "<IMG src = "URL">):

```
<BODY>
<P>I just returned from vacation! Here's a photo of my family at the lake:
<IMG src="http://www.somecompany.com/People/Ian/vacation/family.png"
     alt="A photo of my family at the lake.">
</BODY>
```

This automatic object inclusion may also be achieved with the OBJECT element as follows (an embedded object is designated by "<OBJECT src = "URL">):

```
<BODY>
<P>I just returned from vacation! Here's a photo of my family at the lake:
<OBJECT data="http://www.somecompany.com/People/Ian/vacation/family.png"
        type="image/png">
A photo of my family at the lake.
</OBJECT>
</BODY>
```

The only modification shown by Graber is to the path associated with a link URL. There is nothing in Graber to teach or in any way suggest modifying any embedded object URLs, let alone modifying a URL to refer to a different computer (or domain). And it would not have been obvious to modify Graber to apply to embedded object URLs, since Graber is modifying URLs in order to maintain a path history, and there is no reason to maintain any path history for embedded objects.

Thus, Graber lacks any teaching or suggestion of the claimed "routine for modifying at least one embedded object URL of a web page to include a hostname prepended to a domain name and path." Graber neither teaches nor in any way suggests modifying embedded object URLs, *only links*. And, since Graber only *modifies links* in order to retain path history information and not to go to other servers, Graber neither teaches nor in any way suggests prepending a hostname to a domain name and path.

The Examiner is respectfully reminded that, when not defined by applicant in the specification, the words of a claim must be given their plain meaning. In other words, they must be read as they would be interpreted by those of ordinary skill in the art. *In re Sneed,* 710 F.2d 1544, 218 USPQ 385 (Fed. Cir. 1983). One of ordinary skill in the art would know that an object is not the same as a link.

Claim 41 also recites "a set of repeater servers, distinct from the first server, for hosting at least some of the embedded objects of web pages that are normally hosted by the first server." The Examiner asserts that Graber '979 teaches the set of repeater servers by having "a second server, e.g., OLS *site*, distinct from the first *server*, e.g., 122a, for hosting some of the embedded objects of web pages (cols 10-11)" *Paper No. 23*, pg. 10, item 7(b). The Examiner further states:

> It would have been obvious . . . that the claimed invention differed from the teachings of Graber et al only by a degree, e.g., in the wording of a set of repeater servers by from a broad interpretation of the claims, even taught Grabber [sic] et al did not say that OLS servers were repeater servers, examiner interprets them as equivalent to the repeater servers because they store some of the web pages and serve them to the user." *Paper No. 23*, pgs. 10-11, item 8.

Applicants respectfully disagree. First, there is, in general, a distinction between *sites* and *servers*. A number of sites may be located on the same server. Second, even if, in Graber's example, sites 122a (the Co-marketer #1) and 142 (the OLS Web Server) are on distinct servers, there is nothing in Graber to teach or in any way suggest that, as required by claim 41, site 142 hosts "at least some of the embedded objects of web pages that are normally hosted by" site 122a (or vice versa). The only association between the two sites is that site 122a may have a link URL to the on-line service's home page 128. The mere fact that Graber shows two distinct servers does not in any way support an assertion that Graber anticipates or renders obvious the claimed invention. Graber is completely silent as to any embedded objects, and since the co-marketer sites 122a, 122b and 122c are merely conduits to the on-line service 140, there is no reason why the OLS web server 142 (or any other web server) should host embedded objects of web pages that are normally served by one of the co-marketer sites 122a, etc. Similarly, there is no teaching or suggestion of any kind in Graber of embedded objects normally hosted by OLS web server 142 being served from any other servers.

Graber lacks any teaching or suggestion of the claimed "set of repeater servers, <u>distinct</u> from the first server, for hosting at least some of the embedded objects of web pages that are normally hosted by the first server." The OLS web server is not a distinct server that hosts embedded objects of a web page that are normally hosted by any other server.

Claim 41 further recites that "in response to requests for the web page, generated by the client machines the web page including the modified embedded object URL is served from the first server and the embedded object identified by the modified embedded object URL is served from a given one of the repeater servers."

The Examiner's position (set forth in Paper No. 23 on pg. 10, item 7) is that the co-marketer #1's web page (site) 122a is the first server and that the OLS site 128 is the second server (the repeater server network). The Examiner then states that "in response to requests for

the web page generated by the client machines, the web page including the modified embedded object URL is served from the first server [122a] . . . and the embedded object identified by the modified embedded object URL is served from a given one of the second servers [OLS site 128]." *Paper No. 23,* pg. 10, item 7(c). With respect, this is incorrect. It is not what Graber teaches. The teaching in Graber is that when a user selects (e.g., clicks on) a link on a co-marketer page (e.g., 122a), the user's browser then loads the web page associated with that link. If that link refers to the on-line service (e.g., "WWW.OLS.COM\CM1\INDEX.HTML"), then the user's computer connects to the server denoted "WWW.OLS.COM" and loads the selected page from that server. I.e., in the example given, the user's browser loads the selected page from the OLS Web Server 142. And then all embedded objects associated with that page are also loaded from that server.

In support of his argument, the Examiner refers to Graber's alternate embodiment described in Graber '979 at col. 12, line 35 *et seq.* and col. 13, lines 1 *et seq.* Graber's alternate embodiment tries to achieve the same result as the first embodiment, specifically, maintaining path information so that the system can tell where a user came from (and get the user back there).

> the system insures that the UNIX symbolic link information originally passed to site 128 by a previous web site will be available when OLS 140 attempts to enroll the user into OLS 140. In addition, by inserting the UNIX symbolic link information associated with both a previous web site 122a, b, c and OLS site 140 into page links associated with different web sites (other than site 128), the system permits the user to carry UNIX symbolic link information representing previous location(s) traversed during a user session to further web sites.
> *Graber '979*, col. 15, lines 3-12, emphasis added.

> In this alternate embodiment, the URL used to direct a user from a previous site (e.g., 122a, 122b, 122c) to OLS site 128 includes a string which functions to call a special page_link.cgi program which runs on web server 142. The string passed to OLS site 128 also contains (i) a destination page identifier (or filename) representing the particular web page at site 128 to which the user has been directed by the previous web site, and (ii) a UNIX symbolic link or CMID code associated with the previous web site. More particularly, the destination page identifier and the UNIX symbol link information/CMID code are included in the string as arguments to the page_link.cgi program. An exemplary URL which invokes the page_link.cgi program and that could be used

> by co-marketer site 122a for directing a web user from a site 122a to the home page of site 128 is shown below:
>
> WWW.OLS.COM\page_link.cgi?index@CM1
>
> The first potion (i.e., WWW.OLS.COM) of this exemplary URL identifies web site 128 as the web site to which the user is being directed. The remaining portion (i.e., page_link.cgi?index@CM1) of the URL represents a call to the page_link.cgi program. The program call includes two arguments, namely, a destination page identifier (i.e., index) representing the particular page at site 128 to which the user has been directed, and a UNIX symbolic link/CMID code (i.e., CM1) representing the identity of the web site 122a that directed the user to site 128.
>
> *Graber '979*, col. 11, line 60 to col 12, line 51.

As with the first embodiment described by Graber, the link provided on the co-marketer page (e.g., 122a) is used to get the user to the on-line service home page 128 (on OLS server 142). From then on, until the user explicitly tries to get back to the co-marketer page, the OLS server delivers the content. And the co-marketer page (or site 122a) does not, in response to requests for the OLS web page generated by the client machines, serve the web page including the modified embedded object URL, as required by the claims.

In summary, as to claim 41 and its dependents, Graber lacks any teaching or suggestion of *any* of the claimed elements. Graber does not teach or in any way suggest modifying any embedded object URLs, let alone to "include a hostname prepended to a domain name and path." Graber has no repeater servers and Graber lacks any teaching or suggestion of the claimed framework "wherein in response to requests for the web page, generated by the client machines the web page including the modified embedded object URL is served from the first server and the embedded object identified by the modified embedded object URL is served from a given one of the repeater servers."

Claims 42 and 45 depend on claim 41 and is patentable for at least the same reasons as given above.

In one aspect, e.g., as recited in claim 48, the present invention is a method of serving a page supported at an origin server, the page comprising a markup language base document having associated therewith a set of embedded objects. At least one embedded object is identified by a URL. The method includes rewriting the URL of an embedded object to generate a modified URL. As noted above, neither Graber nor any of the other prior art teaches rewriting

URLs of embedded objects. Graber teaches modifying links, not objects. And further, none of the prior art, alone or in combination teaches the claimed "rewriting the URL of an embedded object to generate a modified URL, the modified URL including a new hostname prepended to an original hostname, wherein the original hostname is maintained as part of the modified URL for use in retrieving the embedded object whenever a cached copy of the embedded object is not available." It would be clear to one of skill in the art that such a rewriting would cause the URL to be resolved to a server denoted by the new hostname instead of the original hostname. Neither Graber nor any of the other prior art teaches rewriting embedded object URLs to resolve to a different computer. And none of the prior art (taken alone or in combination) teaches or in any way suggests keeping the original hostname in a rewritten embedded object URL so that it can be used to retrieve the embedded object if a cached copy of the object is not available.

As recited in claim 48, the method further includes, "in response to a request to serve the page received at the origin server, serving the page with the modified URL;" and "attempting to serve the embedded object from a second server other than origin server, as identified by the new hostname." Again, the prior art (alone or in any combination) fails to teach or suggest the claimed serving of a page from an origin server and trying to serve an embedded object (whose URL has been rewritten) from a different server. Use of the present invention allows a content provider (operating an origin server) to have embedded object URLs rewritten so that the embedded objects are served from servers distinct from the origin server.

Claim 48 further recites "if the cached copy of the embedded object is not available from the second server, obtaining the embedded object from the origin server." This feature is neither taught nor suggested by the prior art (alone or in any combination).

The Examiner is respectfully reminded of the legal standards for a § 103 rejection. Per the MPEP:

> To establish a *prima facie* case of obviousness, three basic criteria must be met. First, there must be some suggestion or motivation, either in the references themselves or in the knowledge generally available to one of ordinary skill in the art, to modify the reference or to combine reference teachings. Second, there must be a reasonable expectation of success. Finally, the prior art reference (or references when combined) must teach or suggest all the claim limitations. The teaching or suggestion to make the claimed combination and the reasonable expectation of success must both be found in the prior art and not based on applicant's

> disclosure. *In re Vaeck,* 947 F.2d 488, 20 USPQ2d 1438 (Fed. Cir. 1991). * * *
>
> The initial burden is on the examiner to provide some suggestion of the desirability of doing what the inventor has done. "To support the conclusion that the claimed invention is directed to obvious subject matter, either the references must * * * suggest the claimed invention or the examiner must present a convincing line of reasoning as to why the artisan would have found the claimed invention to have been obvious in light of the teachings of the references." *Ex parte Clapp*, 227 USPQ 972, 973 (Bd. Pat. App. & Inter. 1985).
>
> *MPEP § 706.02(j)*

Applicant respectfully submits that the Examiner has not met the burden of a *prima facie* case of obviousness. None of the three required criteria has been met. First, none of the prior art references (Graber or Bonnaure) suggest their combination or modification. Second, even if the Examiner's proposed combination of Graber with supposedly known features of networking were suggested in the prior art, no proposed combination of Graber with the known prior art would have produced a system as claimed (i.e., there was no reasonable expectation of success). And finally, as shown above, Graber (or Graber when combined with Bonnaure or any other prior art) does not teach or suggest all the claim limitations.

In view of the above, withdrawal of the claim rejections under 35 U.S.C. § 103 is respectfully requested.

INFORMATION DISCLOSURE STATEMENT

Applicant is filing herewith an Information Disclosure Statement citing art from a European Search report in a counterpart European application.

CONCLUSION

As mentioned in the Personal Interview of 9/20/2001 and in the subsequent telephone conversation, applicant will withdraw the interference request in this application. Applicant respectfully reminds the Examiner, however, that he may, of his own, determine that an interference should be declared between this application and the '703 patent. Should the Examiner so determine, applicant will cooperate fully to get the interference declared.

Should there be any questions or concerns regarding this application, the Examiner is requested to telephone the undersigned.

Applicant respectfully submits that this application is in condition for allowance and an early Action to that effect is earnestly solicited. Applicant seeks an early allowance and expedited issuance of this application.

Respectfully submitted,

By _____

Brian Siritzky
Reg. No. 37,497
Tel. No.: (703) 905-2185
Fax No.: (703) 905-2500

1600 Tysons Blvd.,
McLean, Virginia 22102
(703) 905-2000
30096951v1

APPENDIX

VERSION WITH MARKINGS TO SHOW CHANGES MADE

41. (Amended) A distributed hosting framework operative in a computer network in which users of client machines connect to a first server, the framework comprising:

a routine for modifying at least one embedded object URL of a web page to include a hostname prepended to a domain name and path;

a set of repeater servers, distinct from the first server, for hosting at least some of the embedded objects of web pages that are normally hosted by the first server[;

at least one first level name server that provides a first level domain name service (DNS) resolution; and

at least one second level name server that provides a second level domain name service (DNS) resolution]; and

a repeater selector mechanism constructed and adapted to identify, for a particular client machine, an appropriate repeater server from the set of repeater servers;

wherein in response to requests for the web page, generated by the client machines, the web page including the modified embedded object URL is served from the first server and the embedded object identified by the modified embedded object URL is served from a given one of the repeater servers as identified by the repeater selector mechanism [first level and second level name servers].

45. (Amended) The hosting framework as described in claim 41 wherein the repeater selector mechanism [first level name server] includes a network map for use in directing a request for the embedded object generated by a client.

48. (Amended) A method of serving a page supported at [a first] an origin server, the page comprising a markup language base document having associated therewith a set of embedded objects, [each] at least one embedded object identified by a URL, the method comprising:

rewriting the URL of an embedded object to generate a modified URL, the modified URL including a new hostname prepended to an original hostname, wherein the original hostname is maintained as part of the modified URL for use in retrieving the embedded object whenever a cached copy of the embedded object is not available;

in response to a request to serve the page received at the [first] origin server, serving the page with the modified URL;

attempting to serve the embedded object from a second server other than [first] origin server, as identified by the new hostname; and

if the cached copy of the embedded object is not available from the second server, [serving] obtaining the embedded object from the [first] origin server.

49. (Amended) A method of serving a page and an associated page object, wherein the page is stored on a first server and copies of the page object are stored on a set of servers distinct from the first server, the method comprising:

(a) modifying a URL for the page object to include a hostname prepended to a content provider-supplied domain name and path;

(b) serving the page from the first server with the modified URL;

(c) responsive to a browser query to resolve the hostname, identifying a given one of the set of servers from which the object may be retrieved; and

(d) returning to the browser an [IP] address of the identified server to enable the browser to attempt to retrieve the object from that server.

52. (Amended) The method as described in claim [49] 51 wherein the identified server is selected from a set of repeater servers based on [step of resolving the ARL comprises:

utilizing] data identifying a requesting user's location and data identifying [then-current Internet traffic conditions] current costs between a group containing the requesting user and servers in the set of repeater servers [to identify the server].

55. (Amended) The method as described in claim 54 wherein an identified server is selected from a set of repeater servers based on data identifying a requesting user's location [the identifying step comprises:

resolving a request to the domain as a function of a requesting user's location].

56. (Amended) The method as described in claim 55 wherein [the identifying step comprises:

resolving a request to the domain as a function of a requesting user's location and then-current Internet traffic conditions] an identified server is selected from a set of repeater servers based data identifying a requesting user's location and on data identifying current costs between a group containing the requesting user and the set of repeater servers.

59. (Amended) The method as described in claim 57 [further including the step of rewriting the embedded object URL as the modifies the page] wherein the modifying of the at least one embedded object URL takes place in response to the request for the page.

60. (Amended) The method as described in claim 57 [wherein the step of resolving the hostname includes] further comprising:

identifying a subset of servers that may be available to serve the embedded object based on a location of the client machine and data identifying current costs between a group containing the requesting client machine and a set of repeater servers [current Internet traffic conditions]; and

identifying the server from the subset of servers.

61. (Amended) A content delivery method, comprising:

distributing a set of page objects across a network of repeater servers managed by a domain other than an origin server domain[, wherein the network of servers are organized into a set of regions];

for a given page normally served from the origin server domain, tagging at least some of the embedded objects of the page so that requests for the objects resolve to the repeater server domain instead of the origin server domain; and

in response to a client request for an embedded object of the page:

[resolving the client request as a function of a location of the client machine making the request and current Internet traffic conditions to identify a given region; and]

returning to the client an [IP] address of a given one of the repeater servers within the [given region] repeater domain that is likely to host the embedded object and that is not overloaded.

62. (Amended) A content delivery method, comprising:

tagging an embedded object in a page to resolve to a second domain other than an origin server domain by prepending data to a URL supplied by the origin server to generate a different resource locator;

serving the page with the different resource locator from the origin server;

resolving the different resource locator to identify a server in the second domain; and

serving the embedded object from the identified server.

63. (Amended) The method as described in claim 62 wherein the identified server is selected from a set of repeater servers based on a function of a requesting user's location and [Internet traffic conditions] on data identifying current costs between a group containing the requesting user and the repeater servers.

67. (Amended) A method for Internet content delivery, comprising:

at an origin server, modifying at least one embedded object URL of a page to include a hostname prepended to a domain name and a path normally used to retrieve the embedded object;

responsive to a request for the page issued from a client, serving the page with the modified embedded object URL to the client from the origin server;

responsive to a request for the embedded object, resolving the hostname to an address of a repeater server, other than the [content provider] origin server, that is likely to host the embedded object; and

attempting to serve the embedded object to the client from the repeater server.

APPENDIX SHOWING SUPPORT FOR CLAIMS

This appendix shows some of the support in the specification for some of the claim elements. These tables are not intended to limit the claims in any way and are not intended to show the only support for the claims in the specification. This appendix is being provided only for the convenience of the Examiner.

Claim No. 41	Support in Application
A distributed hosting framework operative in a computer network in which users of client machines connect to a first server, the framework comprising:	See generally Fig. 1 (e.g., client **106**[1] and origin server **102**[2]) and corresponding description, e.g., at pgs. 6 to 8.
a routine for modifying at least one embedded object URL of a web page to include a hostname prepended to a domain name and path;	See generally Fig. 3, "**B5 Rewrite Resource**" and corresponding description, e.g., at pgs. 15, lines 14-15, ("If the resource is an HTML document then the reflector rewrites the HTML document by modifying resource identifiers (URLs) within it"). "The rewriter identifies directives, such as embedded images and anchors, containing URLs." *Pg. 31, lines 8-9*, emphasis added. Resource rewriting is explained and described in greater detail in the section titled "**Rewriting HTML Resources**" on pgs. 30 and 31. One form of the modified URL is specified at pg. 14, lines 13-16, which states: Given a repeater name, scheme, origin server name and path, create a new URL. If the scheme is "http", the preferred embodiment uses the following format: *http://<repeater>/<server>/<path>* See also, e.g., "create a single URL containing the URL of the original resource, as well as the identity of the selected repeater. A special form of URL is created to provide this information." *Pg. 14, lines 8-10.*

[1] "Client **106** is a processor requesting resources from origin server **102** on behalf of an end user. The client **106** is typically a user agent (e.g., a Web browser such as Netscape Communications Corporation's Navigator™) or a proxy for a user agent." *Pg. 7, lines 11-13.*

[2] "Origin server **102** is a server at which resources originate. More generally, the origin server **102** is any process or collection of processes that provide resources in response to requests from a client **106**. Origin server **102** can be any off-the-shelf Web server. In a preferred embodiment, origin server **102** is typically a Web server such as the Apache server or Netscape Communications Corporation's Enterprise™ server." *Pg. 7, lines 6-10.*

Claim No. 41	Support in Application
a set of repeater servers, distinct from the first server, for hosting at least some of the embedded objects of web pages that are normally hosted by the first server;	See generally Fig. 1, repeaters **104**. "Each repeater **104a**, **104b**, and **104c** replicates some or all of the information available on the origin server **102** as well as information available on other origin servers in the network **100**." *Pg. 6, lines 16-18.*
a repeater selector mechanism constructed and adapted to identify, for a particular client machine, an appropriate repeater server from the set of repeater servers	See generally Figs. 1 & 3, reflector **108**, and corresponding description. "the reflector determines (B3-1) the best repeater to reflect the request to" *Pg. 14, lines 3-4.* " * * * then select the best repeater to handle that request * * * This selection is performed by the Best Repeater Selector (BRS) mechanism * * * [which selects], quickly and heuristically, an appropriate repeater for a given client given only the network address of the client. An appropriate repeater is one which is not too heavily loaded and which is not too far from the client in terms of some measure of network distance." *Pg. 19, line 20 to pg. 20, line 1.*
wherein in response to requests for the web page, generated by the client machines the web page including the modified embedded object URL is served from the first server	See generally Figs. 2 & 3, and corresponding description. The browser extracts the host name (origin server/content provider server) from the URL and uses a domain name system (DNS) server to look up a network (IP) address of the corresponding server[3], i.e., to resolve the host name to identify a server hosting that page. The user's browser then uses that IP address to connect to the content provider server[4] and sends a request over that connection for identifying the requested web page.[5] The reflector (possibly at the origin server) analyzes the user's request and determines whether or not to reflect the request[6].or to serve the request from the origin server[7].

---

[3] "A2. The browser extracts the host (origin server) name from the resource identifier, and uses a domain name server (DNS) to look up the network (IP) address of the corresponding server." *The '598 Patent,* col. 7, lines 5-9.

[4] "A3. The browser uses the server's network address and port number to establish a connection between the client **106** and the host or origin server **102**." *The '598 Patent,* col. 7, lines 11-13.

[5] "A4. The client **106** then sends a (GET) request over the connection identifying the requested resource." *The '598 Patent*, col. 7, lines 14-15.

[6] "B1 The reflector **108** analyzes the request to determine whether or not to reflect the request . . . B2 If the request is not from a repeater, . . . the reflector either reflects the request . . . or serves the request locally" *The '598 Patent*, col. 7, line 56 to col. 8, line 6.

Claim No. 41	Support in Application
and the embedded object identified by the modified embedded object URL is served from a given one of the repeater servers as identified by the repeater selector mechanism.	See generally Fig. 2, "A7 Send Reply with Resource" and corresponding description. "The repeater then constructs a reply including the requested resource (which was retrieved from the cache or from the origin server) and sends that reply to the requesting client." *Pg. 19, lines 9-11.*

(... continued)

7 "B4. To serve a request locally, the request is sent by the reflector to ("forwarded to") the origin server **102**. . . . The origin server **102** processes the request in the normal manner (A5-A7)." *The '598 Patent*, col. 8, lines 54-58.

Claim 42	Support in Application
The hosting framework as described in claim 41 wherein a given one of the set of repeater servers includes a buddy server for assuming the hosting responsibilities of the given one of the set of repeater servers upon a given failure condition.	See generally Fig. 1, repeaters **104**. "Each repeater **104a**, **104b**, and **104c** replicates some or all of the information available on the origin server **102** as well as information available on other origin servers in the network **100**." *Pg. 6, lines 16-18.* See also section titled "Repeater Network Resilience" at pg.. 39, line 3 to pg. 41, line 4. "If a master fails ... another repeater will take over the role of master" *Pg. 39, lines 9-10.*

Claim 45	Support in Application
The hosting framework as described in claim 41 wherein a network map is used for directing a request for the embedded object generated by a client.	"The BRS relies on three pre-computed tables, namely the Group Reduction Table, the Link Cost Table, and the Load Table. These three tables (described below) are computed off-line and downloaded to each reflector by its contact in the repeater network" *pg. 20, lines 4-7.* "The Group Reduction Table and Link Cost Table used in BRS processing are created and regularly updated by an independent procedure referred to herein as *NetMap*." *Pg. 24, lines 2-5.* "The *NetMap* procedure ... [creates a] database ... [that] contains essential information used for further processing, namely (1) the identity of each group, (2) the set of IP addresses in a given group, (3) the presence of links between groups indicating paths over which information may travel, and (4) the cost of sending data over a given link." *Pg. 24, line 20 to pg. 25, line 1.* **Location**: "An appropriate repeater is one which is not too heavily loaded and which is not too far from the client in terms of some measure of network distance." *Pg. 19, line 24 to pg. 20, line 25.* "The goal of the BRS is to select * * * an appropriate repeater for a given client given only the network address of the client." *Pg. 19, lines 23-24.* **Conditions**: "The BRS relies on three pre-computed tables, namely the Group Reduction Table, the Link Cost Table, and the Load Table. * * * **The Group Reduction Table** places every network address into a group, with the goal that addresses in a group share relative costs, so that they would have the same best repeater under varying conditions (i.e., the BRS is invariant over the members of the group). . **The Link Cost Table** * * * specifies the current cost between each repeater and each group. * * *Over time, the table will be updated with measurements which more accurately reflect the relative cost of transmitting a file between the repeater and a member of the group." *Pg. 20, lines 4-17.* The BRS computes the cost of servicing a given client from each eligible repeater." *Pg. 22, lines 11-12.* **"Best Repeater Selector (BRS)** * * * Given a client network address and the three tables * * *:

Claim 45	Support in Application
	E1.Determine which group the client is in using the Group Reduction Table. E2. For each repeater in the Link Cost Table and Load Table, determine that repeater's combined cost as follows: * * * E2b. Determine the link cost between the repeater and the client's group (using the Link Cost Table). E2c. Determine the combined cost * * *. E3. Select a small set of repeaters with the lowest cost. E4. Select a random member from the set." *Pg. 23, lines 4-21.*

Claim 48	Support in Application
A method of serving a page supported at an origin server,	See generally Fig. 1 and its corresponding description, e.g., at page 6, line 13 to page 7, line 10. "Origin server **102** is a server at which resources originate. More generally, the origin server **102** is any process or collection of processes that provide resources in response to requests from a client **106**." *Pg. 7, lines 6-8.*
the page comprising a markup language base document having associated therewith a set of embedded objects, at least one embedded object identified by a URL, the method comprising:	"In the case of the Internet in general and the World Wide Web specifically, documents can be created using a standardized form known as the Hypertext Markup Language (HTML). In HTML, a document consists of data (text, images, sounds, and the like), including links to other sections of the same document or to other documents. The links are generally provided as URLs, and can be in relative or absolute form." *Pg. 9, lines 10-16.*
rewriting the URL of an embedded object to generate a modified URL, the modified URL including a new hostname prepended to an original hostname,;	See generally the section titled "**Rewriting HTML Resources**" at pgs.. 30 and 31. All, generally Fig. 3, "B5 Rewrite Resource" and corresponding description, e.g., at pgs. 15, lines 14-15, ("If the resource is an HTML document then the reflector rewrites the HTML document by modifying resource identifiers (URLs) within it"). One form of the modified URL is specified at pg. 14, lines 13-16, which states: D1. Given a repeater name, scheme, origin server name and path, create a new URL. If the scheme is "http", the preferred embodiment uses the following format: *http://<repeater>/<server>/<path>* See also, e.g., "create a single URL containing the URL of the original resource, as well as the identity of the selected repeater. A special form of URL is created to provide this information." *Pg. 14, lines 8-10.*
wherein the original hostname is maintained as part of the modified URL for use in retrieving the embedded object whenever a cached copy of the embedded object is not available	"The repeater then constructs a reply including the requested resource (which was retrieved from the cache or from the origin server) and sends that reply to the requesting client." *Pg. 19, lines 9-11*, emphasis added.

Claim 48	Support in Application
in response to a request to serve the page received at the origin server, serving the page with the modified URL;	**Fig. 5**, "Send GET Request to Repeater" and corresponding description. The browser extracts the host name (origin server/content provider server) from the URL and uses a domain name system (DNS) server to look up a network (IP) address of the corresponding server[8], i.e., to resolve the host name to identify a server hosting that page. The user's browser then uses that IP address to connect to the content provider server[9] and sends a request over that connection for identifying the requested web page.[10] The reflector (possibly at the origin server) analyzes the user's request and determines whether or not to reflect the request[11].or to serve the request from the origin server[12].
Attempting to serve the embedded object from a second server other than the origin server as identified by the new hostname; and	"The repeater then determines whether the requested resource is cached locally. If the requested resource is in the repeater's cache it is retrieved. On the other hand, if a valid copy of the requested resource is not in the repeater's cache, the repeater modifies the incoming URL, creating a request that it issues directly to the originating reflector which processes it (as in B1-B6)." *Pg. 18, line 20 to pg. 19, line 7.*
if the cached copy of the embedded object is not available from the second server, obtaining the embedded object from the origin server.	"If a resource is not cached locally, the cache can query its "peer caches" to see if one of them contains the resource, before or at the same time as requesting the resource from the reflector/origin server." *Pg. 19, lines 2-8.*

---

8 "A2. The browser extracts the host (origin server) name from the resource identifier, and uses a domain name server (DNS) to look up the network (IP) address of the corresponding server." *The '598 Patent,* col. 7, lines 5-9.

9 "A3. The browser uses the server's network address and port number to establish a connection between the client **106** and the host or origin server **102**." *The '598 Patent,* col. 7, lines 11-13.

10 "A4. The client **106** then sends a (GET) request over the connection identifying the requested resource." *The '598 Patent*, col. 7, lines 14-15.

11 "B1 The reflector **108** analyzes the request to determine whether or not to reflect the request . . . B2 If the request is not from a repeater, . . . the reflector either reflects the request . . . or serves the request locally" *The '598 Patent*, col. 7, line 56 to col. 8, line 6.

12 "B4. To serve a request locally, the request is sent by the reflector to ("forwarded to") the origin server **102**. . . . The origin server **102** processes the request in the normal manner (A5-A7)." *The '598 Patent*, col. 8, lines 54-58.

Claim 49	Support in Application
A method of serving a page and an associated page object, wherein the page is stored on a first server and copies of the page object are stored on a set of servers distinct from the first server, comprising	See generally Fig. 1, repeaters **104**. "Each repeater **104a**, **104b**, and **104c** replicates some or all of the information available on the origin server **102** as well as information available on other origin servers in the network **100**." *Pg. 6, lines 16-18.*
(a) modifying a URL for the page object to include a hostname prepended to a content provider-supplied domain name and path;	See generally the section titled "**Rewriting HTML Resources**" at pgs.. 30 and 31. All, generally Fig. 3, "B5 Rewrite Resource" and corresponding description, e.g., at pgs. 15, lines 14-15, ("If the resource is an HTML document then the reflector rewrites the HTML document by modifying resource identifiers (URLs) within it"). One form of the modified URL is specified at pg. 14, lines 13-16, which states: D1. Given a repeater name, scheme, origin server name and path, create a new URL. If the scheme is "http", the preferred embodiment uses the following format: *http://<repeater>/<server>/<path>* See also, e.g., "create a single URL containing the URL of the original resource, as well as the identity of the selected repeater. A special form of URL is created to provide this information." *Pg. 14, lines 8-10.*
(b) serving the page from the first server with the modified URL;	The reflector (possibly at the origin server) analyzes the user's request and determines whether or not to reflect the request[13] or to serve the request from the origin server[14].
(c) responsive to a browser query to resolve the hostname, identifying a given one of the set of servers from which the object may be retrieved; and	**Fig. 5**, "Send GET Request to Repeater" and corresponding description. The browser extracts the host name (origin server/content provider server) from the URL and uses a domain name system (DNS) server to look up a network (IP) address of the corresponding server[15],

---

[13] "B1 The reflector **108** analyzes the request to determine whether or not to reflect the request . . . B2 If the request is not from a repeater, . . . the reflector either reflects the request . . . or serves the request locally" *The '598 Patent*, col. 7, line 56 to col. 8, line 6.

[14] "B4. To serve a request locally, the request is sent by the reflector to ("forwarded to") the origin server **102**. . . . The origin server **102** processes the request in the normal manner (A5-A7)." *The '598 Patent*, col. 8, lines 54-58.

[15] "A2. The browser extracts the host (origin server) name from the resource identifier, and uses a domain name server (DNS) to look up the network (IP) address of the corresponding server." *The '598 Patent,* col. 7, lines 5-9.

Claim 49	Support in Application
	i.e., to resolve the host name to identify a server hosting that page. The user's browser then uses that IP address to connect to the content provider server[16] and sends a request over that connection for identifying the requested web page.[17] The reflector (possibly at the origin server) analyzes the user's request and determines whether or not to reflect the request[18].or to serve the request from the origin server[19]. "The selection . . . of an appropriate repeater to handle the request can be done in a number of ways. . . . from a plurality of repeaters in the network, a set of repeaters is selected, the members of the set having a low cost relative to the cost group which the client is in. ... Then one member of the set is selected ... as the best repeater." *Pg. 4, line 27 to pg. 5, line 6.* See also **Fig. 6, "E3 Select small set of repeaters"** and **"E4 Select Element from set"** and corresponding description at pg. 23.
(d) returning to the browser an address of the identified server to enable the browser to attempt to retrieve the object from that server.	The browser tries automatically to load the embedded object. The DNS returns the IP address corresponding to the URL.

Claim 50	Support in Application
The method as described in claim 15 wherein the copies of the page object are stored on a subset of the set of content servers.	"Each repeater **104a**, **104b**, and **104c** replicates some or all of the information available on the origin server **102** as well as information available on other origin servers in the network **100**." *Pg. 6, lines 16-18*, emphasis added.

(... continued)

16 "A3. The browser uses the server's network address and port number to establish a connection between the client **106** and the host or origin server **102**." *The '598 Patent,* col. 7, lines 11-13.

17 "A4. The client **106** then sends a (GET) request over the connection identifying the requested resource." *The '598 Patent*, col. 7, lines 14-15.

18 "B1 The reflector **108** analyzes the request to determine whether or not to reflect the request . . . B2 If the request is not from a repeater, . . . the reflector either reflects the request . . . or serves the request locally" *The '598 Patent*, col. 7, line 56 to col. 8, line 6.

19 "B4. To serve a request locally, the request is sent by the reflector to ("forwarded to") the origin server **102**. . . . The origin server **102** processes the request in the normal manner (A5-A7)." *The '598 Patent*, col. 8, lines 54-58.

Claim 51	Support in Application
A content delivery method, comprising:	
tagging an embedded object in a page to resolve to a domain other than a content provider domain by prepending given data to a content provider-supplied URL to generate an alternate resource locator (ARL);	"A hypertext document may contain any number of links to other documents, and each of those other documents may be on a different server in a different part of the world." *pg. 9, lines 20-23.* See generally the section titled "**Rewriting HTML Resources**" at pgs.. 30 and 31. All, generally Fig. 3, "**B5 Rewrite Resource**" and corresponding description, e.g., at pgs. 15, lines 14-15, ("If the resource is an HTML document then the reflector rewrites the HTML document by modifying resource identifiers (URLs) within it"). One form of the modified URL is specified at pg. 14, lines 13-16, which states: D1. Given a repeater name, scheme, origin server name and path, create a new URL. If the scheme is "http", the preferred embodiment uses the following format: *http://<repeater>/<server>/<path>* See also, e.g., "create a single URL containing the URL of the original resource, as well as the identity of the selected repeater. A special form of URL is created to provide this information." *Pg. 14, lines 8-10*
serving the page from a content provider server with the ARL; and	**Fig. 5**, "Send GET Request to Repeater" and corresponding description. The browser extracts the host name (origin server/content provider server) from the URL and uses a domain name system (DNS) server to look up a network (IP) address of the corresponding server[20], i.e., to resolve the host name to identify a server hosting that page. The user's browser then uses that IP address to connect to the content provider server[21] and sends a request over that connection for identifying the requested web page.[22]

---

20 "A2. The browser extracts the host (origin server) name from the resource identifier, and uses a domain name server (DNS) to look up the network (IP) address of the corresponding server." *The '598 Patent,* col. 7, lines 5-9.

21 "A3. The browser uses the server's network address and port number to establish a connection between the client **106** and the host or origin server **102**." *The '598 Patent,* col. 7, lines 11-13.

22 "A4. The client **106** then sends a (GET) request over the connection identifying the requested resource." *The '598 Patent*, col. 7, lines 14-15.

Claim 51	Support in Application
	The reflector (possibly at the origin server) analyzes the user's request and determines whether or not to reflect the request[23] or to serve the request from the origin server[24].
resolving the ARL to identify a content server in the domain; and	The browser tries automatically to load the embedded object. The DNS returns the IP address corresponding to the ARL.
serving the embedded object from the identified content server.	"The repeater then constructs a reply including the requested resource (which was retrieved from the cache or from the origin server) and sends that reply to the requesting client." *Pg. 19, lines 9-11.*

Claim 52	Support in Application
The method as described in claim 51 wherein the identified server is selected from a set of repeater servers based on data identifying a requesting user's location and data identifying current costs between a group containing the requesting user and servers in the set of repeater servers .	**Location**: "An appropriate repeater is one which is not too heavily loaded and which is not too far from the client in terms of some measure of network distance." *Pg. 19, line 24 to pg. 20, line 25.* "The goal of the BRS is to select * * * an appropriate repeater for a given client given only the network address of the client." *Pg. 19, lines 23-24.* **Conditions**: "The BRS relies on three pre-computed tables, namely the Group Reduction Table, the Link Cost Table, and the Load Table. * * * **The Group Reduction Table** places every network address into a group, with the goal that addresses in a group share relative costs, so that they would have the same best repeater under varying conditions (i.e., the BRS is invariant over the members of the group). . **The Link Cost Table** * * * specifies the current cost between each repeater and each group. * * *Over time, the table will be updated with measurements which more accurately reflect the relative cost of transmitting a file between the repeater and a member of the group." *Pg. 20, lines 4-17.* 'The BRS computes the cost of servicing a given client from each eligible repeater." *Pg. 22, lines 11-12.* **"Best Repeater Selector (BRS)** * * *

(... continued)

[23] "B1 The reflector **108** analyzes the request to determine whether or not to reflect the request . . . B2 If the request is not from a repeater, . . . the reflector either reflects the request . . . or serves the request locally" *The '598 Patent*, col. 7, line 56 to col. 8, line 6.

[24] "B4. To serve a request locally, the request is sent by the reflector to ("forwarded to") the origin server **102**. . . . The origin server **102** processes the request in the normal manner (A5-A7)." *The '598 Patent*, col. 8, lines 54-58.

Claim 52	Support in Application
	Given a client network address and the three tables * * *: E1.Determine which group the client is in using the Group Reduction Table. E2. For each repeater in the Link Cost Table and Load Table, determine that repeater's combined cost as follows: * * * E2b. Determine the link cost between the repeater and the client's group (using the Link Cost Table). E2c. Determine the combined cost * * *. E3. Select a small set of repeaters with the lowest cost. E4. Select a random member from the set." *Pg. 23, lines 4-21.*

Claim 53	Support in Application
A content delivery service, comprising:	
replicating a set of page objects across a wide area network of servers managed by a domain other than an origin server domain;	See generally Fig. 1, repeaters **104**. "Each repeater **104a**, **104b**, and **104c** replicates some or all of the information available on the origin server **102** as well as information available on other origin servers in the network **100**." *Pg. 6, lines 16-18.*
for a given page normally served from the origin server domain, tagging the embedded objects of the page so that requests for the page objects resolve to the domain instead of the origin server domain;	"See generally the section titled "**Rewriting HTML Resources**" at pgs.. 30 and 31. All, generally Fig. 3, "**B5 Rewrite Resource**" and corresponding description, e.g., at pgs. 15, lines 14-15, ("If the resource is an HTML document then the reflector rewrites the HTML document by modifying resource identifiers (URLs) within it"). "Rewriting requires that a repeater has been selected (as described above with reference to the Best Repeater Selector)." *Pg. 30, lines 14-15.* "For each URL encountered in the resource to be re-written, ... if the URL is repeatable, it is modified to refer to the selected repeater." *Pg. 31, lines 10-14.* One form of the modified URL is specified at pg. 14, lines 13-16, which states: D1. Given a repeater name, scheme, origin server name and path, create a new URL. If the scheme is "http", the preferred embodiment uses the following format: *http://<repeater>/<server>/<path>* See also, e.g., "create a single URL containing the URL of the original resource, as well as the identity of the selected repeater. A special form of URL is created to provide this information." *Pg. 14, lines 8-10*

Claim 53	Support in Application
responsive to a request for the given page received at the origin server domain, serving the given page from the origin server domain; and	**Fig. 5**, "Send GET Request to Repeater" and corresponding description. The browser extracts the host name (origin server/content provider server) from the URL and uses a domain name system (DNS) server to look up a network (IP) address of the corresponding server[25], i.e., to resolve the host name to identify a server hosting that page. The user's browser then uses that IP address to connect to the content provider server[26] and sends a request over that connection for identifying the requested web page.[27] The reflector (possibly at the origin server) analyzes the user's request and determines whether or not to reflect the request[28] or to serve the request from the origin server[29].
serving at least one embedded object of the given page from a given server in the domain instead of from the origin server domain.	The user (client) receives the page from origin server with at least one embedded object's URL rewritten so that it resolves to a different server, possibly in a different domain. Upon receipt of the modified (rewritten) home page, the client's browser then begins to load all the embedded objects in the page. When the browser gets to the resource with the modified URL, with reference to Figure 2, it extracts the server name from the modified URL and resolves that server name to identify a server (in this case one of the repeaters) (at A2) that can handle the request. Then the client/user, establishes a connection with the identified server (at A3) and sends a request for the embedded resource to that server. The repeater then serves the resource to the client.

---

25 "A2. The browser extracts the host (origin server) name from the resource identifier, and uses a domain name server (DNS) to look up the network (IP) address of the corresponding server." *The '598 Patent,* col. 7, lines 5-9.

26 "A3. The browser uses the server's network address and port number to establish a connection between the client **106** and the host or origin server **102**." *The '598 Patent,* col. 7, lines 11-13.

27 "A4. The client **106** then sends a (GET) request over the connection identifying the requested resource." *The '598 Patent*, col. 7, lines 14-15.

28 "B1 The reflector **108** analyzes the request to determine whether or not to reflect the request . . . B2 If the request is not from a repeater, . . . the reflector either reflects the request . . . or serves the request locally" *The '598 Patent*, col. 7, line 56 to col. 8, line 6.

29 "B4. To serve a request locally, the request is sent by the reflector to ("forwarded to") the origin server **102**. . . . The origin server **102** processes the request in the normal manner (A5-A7)." *The '598 Patent*, col. 8, lines 54-58.

Claim 54	Support in Application
The content delivery method as described in claim 51 wherein the serving step comprises:	
for each embedded object, identifying one or more servers from which the embedded object may be retrieved.	"If the resource is an HTML document then the reflector rewrites the HTML document by modifying resource identifiers (URLs) within it" at pg. 15, lines 14-15. One form of the modified URL is specified at pg. 14, lines 13-16. Resource rewriting is explained and described in greater detail in the section titled "**Rewriting HTML Resources**" on pgs. 30 and 31. "Rewriting requires that a repeater has been selected (as described above with reference to the Best Repeater Selector)." *Pg. 30, lines 14-15.* "For each URL encountered in the resource to be re-written, ... if the URL is repeatable, it is modified to refer to the selected repeater." *Pg. 31, lines 10-14.*

Claim 55	Support in Application
The method as described in claim 54 wherein an identified server is selected from a set of repeater servers based on data identifying a requesting user's location .	The goal of the BRS is to select * * * an appropriate repeater for a given client given only the network address of the client." *Pg. 19, lines 23-24.* "An appropriate repeater is one which is not too heavily loaded and which is not too far from the client in terms of some measure of network distance." *Pg. 19, line 24 to pg. 20, line 25.*

Claim 56	Support in Application
The method as described in claim 55 wherein an identified server is selected from a set of repeater servers based data identifying a requesting user's location and on data identifying current costs between a group containing the requesting user and the set of repeater servers.	**Location**: "An appropriate repeater is one which is not too heavily loaded and which is not too far from the client in terms of some measure of network distance." *Pg. 19, line 24 to pg. 20, line 25.* "The goal of the BRS is to select * * * an appropriate repeater for a given client given only the network address of the client." *Pg. 19, lines 23-24.* **Conditions**: "The BRS relies on three pre-computed tables, namely the Group Reduction Table, the Link Cost Table, and the Load Table. * * * **The Group Reduction Table** places every network address into a group, with the goal that addresses in a group share relative costs, so that they would have the same best repeater under varying conditions (i.e., the BRS is invariant over the members of the group). . **The Link Cost Table** * * * specifies the current cost between each repeater and each group. * * *Over time, the table will be updated with measurements which more accurately reflect the relative cost of transmitting a file between the repeater and a member of the group." *Pg. 20, lines 4-17.* 'The BRS computes the cost of servicing a given client from each eligible repeater." *Pg. 22, lines 11-12.* **"Best Repeater Selector (BRS)** * * * Given a client network address and the three tables * * *: E1.Determine which group the client is in using the Group Reduction Table. E2. For each repeater in the Link Cost Table and Load Table, determine that repeater's combined cost as follows: * * * E2b. Determine the link cost between the repeater and the client's group (using the Link Cost Table). E2c. Determine the combined cost * * *. E3. Select a small set of repeaters with the lowest cost. E4. Select a random member from the set." *Pg. 23, lines 4-21.*

Claim 57	Support in Application
A method for Internet content delivery, comprising:	
at a first server, modifying at least one embedded object URL of a page to include a hostname prepended to a domain name and a path normally used to retrieve the embedded object;	See generally the section titled "**Rewriting HTML Resources**" at pgs.. 30 and 31. Also, generally Fig. 3, "B5 Rewrite Resource" and corresponding description, e.g., at pgs. 15, lines 14-15, ("If the resource is an HTML document then the reflector rewrites the HTML document by modifying resource identifiers (URLs) within it"). One form of the modified URL is specified at pg. 14, lines 13-16, which states:

Claim 57	Support in Application
	D1. Given a repeater name, scheme, origin server name and path, create a new URL. If the scheme is "http", the preferred embodiment uses the following format: *http://<repeater>/<server>/<path>* See also, e.g., "create a single URL containing the URL of the original resource, as well as the identity of the selected repeater. A special form of URL is created to provide this information." *Pg. 14, lines 8-10.*
responsive to a request for the page issued from a client machine, serving the page with the modified embedded object URL to the client machine from the first server;	See generally Fig. 3 and corresponding description at pgs. 14-15. "upon receipt of a request, B1 The reflector ... analyzes the request to determine whether or not to reflect the request. ... B5. If the resource is an HTML document then the reflector rewrites the HTML document by modifying resource identifiers (URLs) within it, ... The resource, possibly as modified by rewriting, is then returned in a reply to the requesting client" Pg. 13, line 2 to pg. 15, line 18. **Fig. 5**, "Send GET Request to Repeater" and corresponding description. The browser extracts the host name (origin server/content provider server) from the URL and uses a domain name system (DNS) server to look up a network (IP) address of the corresponding server[30], i.e., to resolve the host name to identify a server hosting that page. The user's browser then uses that IP address to connect to the content provider server[31] and sends a request over that connection for identifying the requested web page.[32]

30 "A2. The browser extracts the host (origin server) name from the resource identifier, and uses a domain name server (DNS) to look up the network (IP) address of the corresponding server." *The '598 Patent,* col. 7, lines 5-9.

31 "A3. The browser uses the server's network address and port number to establish a connection between the client **106** and the host or origin server **102**." *The '598 Patent,* col. 7, lines 11-13.

32 "A4. The client **106** then sends a (GET) request over the connection identifying the requested resource." *The '598 Patent*, col. 7, lines 14-15.

Claim 57	Support in Application
	The reflector (possibly at the origin server) analyzes the user's request and determines whether or not to reflect the request[33] or to serve the request from the origin server[34].
responsive to a request for the embedded object, resolving the hostname to an IP address of a server other than the first server, that is likely to host the embedded object; and	The embedded object will have a modified URL. Pg. 13, line 2 to pg. 15, line 18.
attempting to serve the embedded object to the client from the server.	"The repeater then constructs a reply including the requested resource (which was retrieved from the cache or from the origin server) and sends that reply to the requesting client." *Pg. 19, lines 9-12.*

Claim 58	Support in Application
The method as described in claim 57 wherein the page is formatted according to a markup language.	"In the case of the Internet in general and the World Wide Web specifically, documents can be created using a standardized form known as the Hypertext Markup Language (HTML)." *Pg. 9, lines 10-12.*

(... continued)

[33] "B1 The reflector **108** analyzes the request to determine whether or not to reflect the request . . . B2 If the request is not from a repeater, . . . the reflector either reflects the request . . . or serves the request locally" *The '598 Patent*, col. 7, line 56 to col. 8, line 6.

[34] "B4. To serve a request locally, the request is sent by the reflector to ("forwarded to") the origin server **102**. . . . The origin server **102** processes the request in the normal manner (A5-A7)." *The '598 Patent*, col. 8, lines 54-58.

Claim 59	Support in Application
The method as described in claim 57 further including the step of rewriting the embedded object URL as the first server modifies the page.	See generally the section titled "**Rewriting HTML Resources**" at pgs.. 30 and 31.

Claim 60	Support in Application
The method as described in claim 23 wherein the step of resolving the hostname includes:	
identifying a subset of servers that may be available to serve the embedded object based on a location of the client machine and data identifying current costs between a group containing the requesting client machine and a set of repeater servers ; and	"The selection ... of an appropriate repeater to handle the request can be done in a number of ways. In the preferred embodiment, it is done by first pre-partitioning the network into "cost groups" and then determining which cost group the client is in. Next, from a plurality of repeaters in the network, a set of repeaters is selected, the members of the set having a low cost relative to the cost group which the client is in." *Pg. 4, line 27 to pg. 5, line 5.* "...requests from certain network addresses (e.g., requests from clients on the same local area network as the reflector itself) are never reflected. Also, the reflector may choose not to reflect requests because the reflector is exceeding its committed aggregate information rate" *Pg. 16, lines 11-14* See also pg. 15, lines 26-28.
identifying the server from the subset of servers.	"Then one member of the set is selected, ..., as the best repeater." *Pg. 5, lines 4-5.*

Claim 61	Support in Application
A content delivery method, comprising:	
distributing a set of page objects across a network of repeater servers managed by a domain other than an origin server domain;	See generally Fig. 1, repeaters **104**. "Each repeater **104a**, **104b**, and **104c** replicates some or all of the information available on the origin server **102** as well as information available on other origin servers in the network **100**." *Pg. 6, lines 16-18.*
for a given page normally served from the origin server domain, tagging at least some of the embedded objects of the page so that requests for the objects resolve to the repeater server domain instead of the origin server domain; and	See generally the section titled "**Rewriting HTML Resources**" at pgs.. 30 and 31. "If the resource is an HTML document then the reflector rewrites the HTML document by modifying resource identifiers (URLs) within it" at pg. 15, lines 14-15. One form of the modified URL is specified at pg. 14, lines 13-16. "Rewriting requires that a repeater has been selected (as described above with reference to the Best Repeater Selector)." *Pg. 30, lines 14-15.* "For each URL encountered in the resource to be re-written, ... if the URL is repeatable, it is modified to refer to the selected repeater." *Pg. 31, lines 10-14.*
in response to a client request for an embedded object of the page:	See generally Fig. 3 and corresponding description at pgs. 14-15.

Claim 61	Support in Application
returning to the client an address of a given one of the repeater servers within the repeater domain that is likely to host the embedded object and that is not overloaded.	The user (client) receives the page from origin server with at least one embedded object's URL rewritten so that it resolves to a different server, possibly in a different domain. Upon receipt of the modified (rewritten) home page, the client's browser then begins to load all the embedded objects in the page. When the browser gets to the resource with the modified URL, with reference to Figure 2, it extracts the server name from the modified URL and resolves that server name to identify a server (in this case one of the repeaters) (at A2) that can handle the request. Then the client/user, establishes a connection with the identified server (at A3) and sends a request for the embedded resource to that server. The repeater then serves the resource to the client. Since the URL for the embedded object refers to a repeater that was selected based on load, when that URL is resolved, it will refer to a repeater that is not overloaded. See, Best Repeater Selector (BRS) at pg. 23, lines 11-14 and Fig. 6. "E2. For each repeater * * * determine that repeater's combined cost as follows: E2a. Determine the maximum and current load on the repeater (using the Load Table)."

Claim 62	Support in Application
A content delivery method, comprising:	
tagging an embedded object in a page to resolve to a second domain other than an origin server domain by prepending data to a URL supplied by the origin server to generate a different resource locator;	See generally the section titled "**Rewriting HTML Resources**" at pgs.. 30 and 31. "If the resource is an HTML document then the reflector rewrites the HTML document by modifying resource identifiers (URLs) within it" at pg. 15, lines 14-15. One form of the modified URL is specified at pg. 14, lines 13-16. "Rewriting requires that a repeater has been selected (as described above with reference to the Best Repeater Selector)." *Pg. 30, lines 14-15.* "For each URL encountered in the resource to be re-written, ... if the URL is repeatable, it is modified to refer to the selected repeater." *Pg. 31, lines 10-14.*
serving the page with the different resource locator from the origin server;	See generally Fig. 3 and corresponding description at pgs. 14-15.
resolving the different resource locator to identify a server in the second domain; and	Pg. 17, lines 1-14, show an origin server in the ".com" domain and a repeater being in the ".edu" domain. "* * * if www.example.com is the origin server, names for three repeaters might be created: wr1.example.com wr2.example.com wr3.example.com The name "wr1.example.com" would be an alias for repeater 1, which might also be known by other names such as "wr1.anotherExample.com" and "wr1.example.edu". If the repeater can determine by which name it was addressed, it can use this information (along with a table that associates repeater alias names with origin server names) to determine which origin server is being addressed. For instance, if repeater 1 is addressed as wr1.example.com, then the origin server is "www.example.com"; if it is addressed as wr1.anotherExample.com", then the origin server is"www.anotherExample.com".

Claim 62	Support in Application
serving the embedded object from the identified server.	"The repeater then constructs a reply including the requested resource (which was retrieved from the cache or from the origin server) and sends that reply to the requesting client." *Pg. 19, lines 9-12.* *See, e.g., support for claims 51, 53, 57*

Claim 63	Support in Application
The method as described in claim 62 wherein the identified server is selected from a set of repeater servers based on a function of a requesting user's location and on data identifying current costs between a group containing the requesting user and the repeater servers.	See claim 56.

Claim 65	Support in Application
A content delivery service, comprising:	
replicating a set of page objects across a network of repeater servers;	See generally Fig. 1, repeaters **104**. "Each repeater **104a**, **104b**, and **104c** replicates some or all of the information available on the origin server **102** as well as information available on other origin servers in the network **100**." *Pg. 6, lines 16-18.*
for a given page normally served from an origin server, tagging at least one embedded object of the page so that requests for the page object resolve to one of the repeater servers instead of to the origin server;	See generally the section titled "**Rewriting HTML Resources**" at pgs.. 30 and 31. "If the resource is an HTML document then the reflector rewrites the HTML document by modifying resource identifiers (URLs) within it" at pg. 15, lines 14-15. One form of the modified URL is specified at pg. 14, lines 13-16. "Rewriting requires that a repeater has been selected (as described above with reference to the Best Repeater Selector)." *Pg. 30, lines 14-15.* "For each URL encountered in the resource to be re-written, ... if the URL is repeatable, it is modified to refer to the selected repeater." *Pg. 31, lines 10-14.*
responsive to a request for the given page received at the origin server, serving the given page from the origin server; and	See generally Fig. 3 and corresponding description at pgs. 14-15.
serving at least one embedded object of the given page from a repeater server instead of from the origin server.	"The repeater then constructs a reply including the requested resource (which was retrieved from the cache or from the origin server) and sends that reply to the requesting client." *Pg. 19, lines 9-12.* *See, e.g., support for claims 51, 53, 57*

Claim 66	Support in Application
66. A service as in claim 65 wherein the origin server and the repeater servers are in different domains.	Pg. 17, lines 1-14, show an origin server in the ".com" domain and a repeater being in the ".edu" domain. "* * * if www.example.com is the origin server, names for three repeaters might be created: wr1.example.com * * * The name "wr1.example.com" would be an alias for repeater 1, which might also be known by other names such as "wr1.anotherExample.com" and "wr1.example.edu". If the repeater can determine by which name it was addressed, it can use this information (along with a table that associates repeater alias names with origin server names) to determine which origin server is being addressed. For instance, if repeater 1 is addressed as wr1.example.com, then the origin server is "www.example.com"; if it is addressed as wr1.anotherExample.com", then the origin server is"www.anotherExample.com".

Claim 67	Support in Application
A method for Internet content delivery, comprising:	
at an origin server, modifying at least one embedded object URL of a page to include a hostname prepended to a domain name and a path normally used to retrieve the embedded object;	See generally the section titled "**Rewriting HTML Resources**" at pgs.. 30 and 31. "If the resource is an HTML document then the reflector rewrites the HTML document by modifying resource identifiers (URLs) within it" at pg. 15, lines 14-15. One form of the modified URL is specified at pg. 14, lines 13-16. "Rewriting requires that a repeater has been selected (as described above with reference to the Best Repeater Selector)." *Pg. 30, lines 14-15.* "For each URL encountered in the resource to be re-written, ... if the URL is repeatable, it is modified to refer to the selected repeater." *Pg. 31, lines 10-14.*
Responsive to a request for the page issued from a client, serving the page with the modified embedded object URL to the client from the origin server;	See generally Fig. 3 and corresponding description at pgs. 14-15.
Responsive to a request for the embedded object, resolving the hostname to an address of a repeater server, other than the origin server, that is likely to host the embedded object; and	See claim 62
Attempting to serve the embedded object to the client from the repeater server.	"The repeater then constructs a reply including the requested resource (which was retrieved from the cache or from the origin server) and sends that reply to the requesting client." *Pg. 19, lines 9-12.* *See, e.g., support for claims 51, 53, 57*

Claim 68	Support in Application
The method as described in claim 67 wherein the page is formatted according to a markup language.	See generally the section titled "**Rewriting HTML Resources**" at pgs.. 30 and 31.

Claim 69	Support in Application
A content delivery service, comprising:	
Replicating a set of page objects across a wide area network of repeater servers;	See generally Fig. 1, repeaters **104**. "Each repeater **104a**, **104b**, and **104c** replicates some or all of the information available on the origin server **102** as well as information available on other origin servers in the network **100**." *Pg. 6, lines 16-18.*
for a given page normally served from an origin server, tagging at least one embedded object of the page so that requests for the page objects resolve to one of the repeater servers instead of to the origin server;	See generally the section titled "**Rewriting HTML Resources**" at pgs.. 30 and 31. "If the resource is an HTML document then the reflector rewrites the HTML document by modifying resource identifiers (URLs) within it" at pg. 15, lines 14-15. One form of the modified URL is specified at pg. 14, lines 13-16. "Rewriting requires that a repeater has been selected (as described above with reference to the Best Repeater Selector)." *Pg. 30, lines 14-15.* "For each URL encountered in the resource to be re-written, ... if the URL is repeatable, it is modified to refer to the selected repeater." *Pg. 31, lines 10-14.*
in response to a request for the given page received at the origin server, causing the given page to be served from the origin server; and	See generally Fig. 3 and corresponding description at pgs. 14-15.
Serving at least one embedded object of the given page from a repeater server instead of from the origin server.	See Figs. 2 and 3 and corresponding description.

Claim 70	Support in Application
70. A method of serving a page and an associated embedded page object, wherein the page is stored on an origin server and copies of the page object are stored on a set of repeater servers distinct from the origin server, wherein each repeater server replicates some or all of the information available on the origin server, the method comprising:	See generally Fig. 1, repeaters **104**. "Each repeater **104a**, **104b**, and **104c** replicates some or all of the information available on the origin server **102** as well as information available on other origin servers in the network **100**." *Pg. 6, lines 16-18.*
(a) causing the embedded page object to resolve to a server in the repeater server domain instead of to the origin server;	See generally the section titled "**Rewriting HTML Resources**" at pgs.. 30 and 31. "If the resource is an HTML document then the reflector rewrites the HTML document by modifying resource identifiers (URLs) within it" at pg. 15, lines 14-15.

Claim 70	Support in Application
	One form of the modified URL is specified at pg. 14, lines 13-16. "Rewriting requires that a repeater has been selected (as described above with reference to the Best Repeater Selector)." *Pg. 30, lines 14-15.* "For each URL encountered in the resource to be re-written, ... if the URL is repeatable, it is modified to refer to the selected repeater." *Pg. 31, lines 10-14.*
(b) serving the page from the origin server; and	See generally Fig. 3 and corresponding description at pgs. 14-15.
(c) responsive to a browser request for the page object, identifying a given repeater server from the set of repeater servers from which the object may be retrieved; and	See, e.g., claim 65.
(d) returning to the browser an address of the identified repeater server to enable the browser to attempt to retrieve the object from that server.	See, e.g., claim 65.

Claim 71	Support in Application
A content delivery method operable in an Internet framework which includes an origin server and at least one set of repeater servers distinct from the origin server, wherein each repeater server replicates some or all of the information available on the origin server, the content delivery method comprising:	See generally Fig. 1, repeaters **104**. "Each repeater **104a**, **104b**, and **104c** replicates some or all of the information available on the origin server **102** as well as information available on other origin servers in the network **100**." *Pg. 6, lines 16-18.*
(A) causing an embedded object in a web page to resolve to a repeater server instead of the origin server; and	See generally the section titled "**Rewriting HTML Resources**" at pgs.. 30 and 31. "If the resource is an HTML document then the reflector rewrites the HTML document by modifying resource identifiers (URLs) within it" at pg. 15, lines 14-15. One form of the modified URL is specified at pg. 14, lines 13-16. "Rewriting requires that a repeater has been selected (as described above with reference to the Best Repeater Selector)." *Pg. 30, lines 14-15.* "For each URL encountered in the resource to be re-written, ... if the URL is repeatable, it is modified to refer to the selected repeater." *Pg. 31, lines 10-14.*
(B) responsive to a request from a client for the web page,	See generally Fig. 3 and corresponding description at pgs. 14-15. "The repeater then constructs a reply including the requested resource (which was retrieved from the cache or from the origin server) and sends that reply to the requesting client." *Pg. 19, lines 9-12.* *See, e.g., support for claims 51, 53, 57*
(a) serving the web page from the origin server; and	
(b) serving the embedded object from the repeater server.	

Claim 72	Support in Application
A method as in claim 71 wherein the repeater server is selected based on data identifying the requesting client's location and on data identifying current costs between a group containing the requesting client and the repeater servers.	See claim 56.

Claim 73	Support in Application
73. A content delivery method, in a system in which a network of repeater servers replicate some or all of the information available on at least one origin server distinct from the repeater servers, the method comprising:	See generally Fig. 1, repeaters **104**. "Each repeater **104a**, **104b**, and **104c** replicates some or all of the information available on the origin server **102** as well as information available on other origin servers in the network **100**." *Pg. 6, lines 16-18.*
Causing an embedded object in a page to resolve to a repeater server instead of an origin server;	See generally the section titled "**Rewriting HTML Resources**" at pgs.. 30 and 31. "If the resource is an HTML document then the reflector rewrites the HTML document by modifying resource identifiers (URLs) within it" at pg. 15, lines 14-15. One form of the modified URL is specified at pg. 14, lines 13-16. "Rewriting requires that a repeater has been selected (as described above with reference to the Best Repeater Selector)." *Pg. 30, lines 14-15.* "For each URL encountered in the resource to be re-written, ... if the URL is repeatable, it is modified to refer to the selected repeater." *Pg. 31, lines 10-14.*
Serving the page from an origin server; Resolving a URL of the embedded object to identify a repeater server; and Serving the embedded object from the identified repeater server.	See generally Fig. 3 and corresponding description at pgs. 14-15. "The repeater then constructs a reply including the requested resource (which was retrieved from the cache or from the origin server) and sends that reply to the requesting client." *Pg. 19, lines 9-12.* *See, e.g., support for claims 51, 53, 57*

Claim 74	Support in Application
An Internet method of serving an HTML page supported at an origin server, the page comprising an HTML base document having associated therewith a set of embedded objects, the method comprising:	See, e.g., claim 48
Causing at least one embedded object to resolve to a repeater server instead of the origin server,	See generally the section titled "**Rewriting HTML Resources**" at pgs.. 30 and 31. "If the resource is an HTML document then the reflector rewrites the HTML document by modifying resource identifiers (URLs) within it" at pg. 15, lines 14-15. One form of the modified URL is specified at pg. 14, lines 13-16. "Rewriting requires that a repeater has been selected (as described above with reference to the Best Repeater Selector)." *Pg. 30, lines 14-15.*

Claim 74	Support in Application
	"For each URL encountered in the resource to be re-written, ... if the URL is repeatable, it is modified to refer to the selected repeater." *Pg. 31, lines 10-14.*
in response to a request to serve the page received at the origin server, serving the page;	See generally Fig. 3 and corresponding description at pgs. 14-15.
Attempting to serve the embedded object from a repeater server; and	See, e.g., claim 48.
if a cached copy of the embedded object is not available from the repeater server, obtaining the embedded object from the origin server.	

Claim 75	Support in Application
A method as in claim 51 wherein requesting client's is a member of a group based on the client's location, and wherein the identified server is selected from a set of repeater servers based on a relative cost of transmitting a file between the repeater and a member of the group.	See claim 56

Claim 76	Support in Application
A method as in claim 48 wherein the rewriting of the URL is in response to a client's request to serve the page.	See generally the section titled "**Rewriting HTML Resources**" at pgs.. 30 and 31. "If the resource is an HTML document then the reflector rewrites the HTML document by modifying resource identifiers (URLs) within it" at pg. 15, lines 14-15. See also Fig. 3.

Claim 77	Support in Application
A method as in claim 49 wherein the modifying of the URL for the page object is performed in response to a request for the page.	See generally the section titled "**Rewriting HTML Resources**" at pgs.. 30 and 31. "If the resource is an HTML document then the reflector rewrites the HTML document by modifying resource identifiers (URLs) within it" at pg. 15, lines 14-15. See also Fig. 3.

Claim 78	Support in Application
78. A method as in claim 51 wherein the tagging of the embedded object in the page takes place in response to a request for the page.	See generally the section titled "**Rewriting HTML Resources**" at pgs. 30 and 31. "If the resource is an HTML document then the reflector rewrites the HTML document by modifying resource identifiers (URLs) within it" at pg. 15, lines 14-15. See also Fig. 3.

Claim 79	Support in Application
79. In a computer network in which client computers request HTML resources from origin servers, and wherein a set of repeater servers, distinct from the origin	See generally Fig. 1, repeaters **104**. "Each repeater **104a**, **104b**, and **104c** replicates some or all of the information available on the origin server **102** as well as information available on other origin servers in the network **100**." *Pg. 6, lines 16-*

Claim 79	Support in Application
server, host at least some of the embedded objects of web pages that are normally hosted by the origin server:	*18.*
a routine constructed and adapted to cause at least one embedded object URL of a web page to be served from one of the repeater servers instead of from the origin server,	See generally the section titled "**Rewriting HTML Resources**" at pgs.. 30 and 31.
Wherein in response to requests for the web page, generated by the client computers, the web page is served from the origin server and the embedded object is served from a given one of the repeater servers.	See generally Fig. 3 and corresponding description at pgs. 14-15.

Claim 81	Support in Application
A method of serving a page supported at an origin server, the page comprising a markup language base document having associated therewith a set of embedded objects, at least one embedded object identified by a URL, the method comprising:	See, e.g., claims 48 and 49.
in response to a request to serve the page received at the origin server:	See, e.g., claims 48 and 49.
(a) rewriting the URL of an embedded object to generate a modified URL, the modified URL including a new hostname prepended to an original hostname, wherein the original hostname is maintained as part of the modified URL for use in retrieving the embedded object whenever a cached copy of the embedded object is not available;	See, e.g., claims 48 and 49.
(b) serving the page with the modified URL;	See, e.g., claims 48 and 49.
(c) attempting to serve the embedded object from a second server other than origin server, as identified by the new hostname; and	See, e.g., claims 48 and 49.
(d) if the cached copy of the embedded object is not available from the second server, attempting to obtain the embedded object from the origin server.	See, e.g., claims 48 and 49.

Claim 82	Support in Application
A content delivery service comprising:	See, e.g., claims 48, 49.

Claim 82	Support in Application
(A) causing a set of page objects to be replicated across a wide area network of servers managed by a domain other than an origin server domain;	
(B) for a given page normally served from the origin server domain, causing the embedded objects of the page to the domain instead of the origin server domain; and	
(C) responsive to a request for the given page received at the origin server domain,	
(C1) serving the given page from the origin server domain; and	
(C2) serving at least one embedded object of the given page from a given server in the domain instead of from the origin server domain.	

Claim 83	Support in Application
A content delivery method comprising:	See, e.g., claims 48-49.
(a) causing a set of page objects to be replicated across a network of repeater servers;	
(b) for a given page that is normally served from an origin server, modifying at least one embedded object of the page so that requests for the page object resolve to one of the repeater servers instead of to the origin server;	
(c) in response to a request for the given page received at the origin server,	
(c1) serving the given page from the origin server; and	
(c2) causing at least one embedded object of the given page to be served from a repeater server instead of from the origin server.	

Claim 84	Support in Application
A framework as in claim 41 wherein the repeater selector mechanism is co located with the first server.	The selector mechanism is preferably located at the reflector 108 which may be co-located at the origin (first) server.

Claim 85	Support in Application
A framework as in claim 41 wherein the repeater selector mechanism is constructed and adapted to identify the appropriate repeater server based on a load on the repeater servers and on a measure of network distance from the client to the various repeater servers.	See generally Figs. 1 & 3, reflector **108**, and corresponding description. "the reflector determines (B3-1) the best repeater to reflect the request to" *Pg. 14, lines 3-4.* " * * * then select the best repeater to handle that request * * * This selection is performed by the Best Repeater Selector (BRS) mechanism * * * [which selects], quickly and heuristically, an appropriate repeater for a given client given only the network address of the client. An appropriate repeater is one which is not too heavily loaded and which is not too far from the client in terms of some measure of network distance." *Pg. 19, line 20 to pg. 20, line 1.*

Claim 86	Support in Application
A distributed hosting framework operative in a computer network in which users of client machines connect to a first server, the framework comprising:	See claim 41
a routine for modifying at least one embedded object URL of a web page to include a hostname prepended to a domain name and path;	
a set of repeater servers, distinct from the first server, for hosting at least some of the embedded objects of web pages that are normally hosted by the first server; and	
a plurality of repeater selector mechanisms, each constructed and adapted to identify, for a particular client machine, a server from a set of repeater servers;	See Claims 41 & 85. Also, there may be a plurality of reflectors, each of which may have a repeater selector mechanism. '**FIGURE 1** shows only a part of a network **100** according to this invention. A complete operating network consists of any number of * * * reflectors * * *." *Pg. 8, lines 20-22.* "* * * the reflector **108** must then select the best repeater to handle that request * * *. This selection is performed by the Best Repeater Selector (BRS) mechanism." *P.g 19, line 20 to pg. 20, line 1.*
wherein in response to requests for the web page, generated by the client machines, the web page including the modified embedded object URL is served from the first server and the embedded object identified by the modified embedded object URL is served from a given one of the repeater servers as identified by at least one of the repeater selector mechanisms.	See Claim 41